arXiv:1604.08402v1 [stat.ML] 28 Apr 2016

Two Differentially Private Rating Collection Mechanisms for Recommender Systems

Zheng Wenjie

April 29, 2016

Abstract

We design two mechanisms for the recommender system to collect user ratings. One is modified Laplace mechanism, and the other is randomized response mechanism. We prove that they are both differentially private and preserve the data utility.

1 Introduction

Recommender Systems (RS) [1] are a kind of system that seek to recommend to users what they are likely interested in. Unlike search engines, the users do not need to type any keyword. The RS's will learn their interest automatically. For instance, if the user has just bought a numeric camera, the RS will recommend to him some SD memory cards; if a user watches a lot of action movies, the RS may suggest some other action movies to him. And this is the typical behaviors which we observe universally in Netflix (movies), Youtube (videos), Google Play (apps), Facebook (friends), Amazon (goods) and other platforms today.

One may wonder how it works. Let us take Netflix as an example. Netflix has a mechanism that allows every user to rate the movies they have watched. Based on these ratings, Netflix builds a profile for each user. And there should be quite a few methods to predict the user preference on other movies that the user has yet seen. Readers can learn more in Section 3. This is not the main topic of this article.

The issue addressed in this article is whether user privacy is compromised by the rating collection mechanism, and what we may do to prevent it. In 2006, Netflix published 100 480 507 ratings that 480 189 users gave to 17 770 movies on the Internet to hold the *Netflix Prize* competition [2]. The data are anonymous. However, in 2007, two researchers from the University of Texas, de-anonymized some of the Netflix data by matching the data set with movie ratings on the Internet Movie Database [3]. This aroused big privacy concern. In 2009, four Netflix users filed a lawsuit against Netflix. We see that this concern of privacy leak is real, and the anonymization alone is not sufficient to prevent it.

Let us take a close look. There were actually two privacy leaks during the procedure. Firstly, the users leaked their ratings to the service provider Netflix. Then, Netflix leaked the ratings to the public. Users made a big fuss on the second leak, but they overlooked the fact that it was themselves who leaked the ratings to Netflix at the first place. Usually, all legal companies will ask the users to sign a user agreement, which authorizes the companies to collect user data and to use them for certain purposes. However, almost no users will ever read it. Anyway, if they do not agree, they will not be able to use the service.

Hence, the goal of this article is to minimize the privacy leak but still guaranteeing the functionality of the service provider. We will achieve this goal by building differential privacy (DP) into the rating collection mechanism. The concept of DP will be explained in detail in Section 2. The big idea is that the user ratings are transformed through the rating collection mechanism, so that from the output (transformed ratings), one cannot know for sure what the input (original ratings) is. Of

course, this kind of transform should satisfy certain properties. After this transform, the service provider can do whatever they want with the ratings without worrying about privacy leak. They can either analyze it themselves, or subcontract the work to a third party by giving them the data access. It is also possible for Netflix to hold a second competition.

One trivial transform is to transform every rating to zero or pure random number. This absolutely prevents any privacy leak, but it erases all information contained in the data as well. Therefore, when we build DP into the mechanism, we should be careful in order to preserve as much information as possible in the data. We designed two mechanisms. One is *modified Laplace mechanism*, and the other is *randomized response mechanism* (Section 2). We will show that they preserve the utility of the ratings (Section 3).

Related work

[4] also tries to bring DP to RS. Their method is different from ours. Let X^i be the original rating set of the i-th user, S be some aggregation statistic of ratings, $\mathcal{A}$ be some algorithm to do data analysis, and f be some transform that guarantees DP. Their method can be summarized as $\mathcal{A}\left(f\left(S\left(\otimes_{i=1}^{n} X^i\right)\right)\right)$, while our method, with a little abuse of notation, can be summarized as $\mathcal{A}\left(S\left(\otimes_{i=1}^{n} f\left(X^i\right)\right)\right)$.

Note that we changed the position of the transform f. This modification is of significant advantage. In their method, f should be adapted to each statistic S, and they can only use algorithms $\mathcal{A}$ relying on S. In our method, we can generalize it to $\mathcal{A}\left(\otimes_{i=1}^{n} f\left(X^i\right)\right)$, which means that we can use more types of algorithms. Furthermore, as long as f is "conjugate" (i.e. f does not change the space where X^i is in), all previous successful algorithms could be seamlessly "transplanted". And we will illustrate in Section 3 that this transplantation is also seamless in theoretical guarantee.

When giving it a second thought, their method is neither privacy preserved nor meaningful. According to their method, at the moment where users transfer their ratings to the service provider so that it could calculate the statistic S, user privacy has already leaked to the service provider. Then, the service provider sends a differentially private version of recommendation back to the user. But why the user bothers to protect the privacy against himself?!

2 Mechanisms

In this section, we will first introduce the concept of *differential privacy*. Then we will define the modified Laplace mechanism and randomized response mechanism. Throughout this section, we consider the rating vector of a single user: $x = (x_1, x_2, \ldots, x_n)$, where n is the number of items. Note that the components may have missing values.

2.1 Differentially private mechanism

The concept of DP is not some entity lurking in the data, but it is used to describe a certain type of data releasing mechanism. The original idea is introduced in [5]. Since, dozens of formulations came out. We use the simplest formulation here.

Definition 1. Let ϵ be a positive value, a random application $\mathcal{M} : \mathfrak{R} \longrightarrow \mathfrak{S}$ is called ϵ-differentially private mechanism if

$$\Pr\left(\mathcal{M}\left(y\right) \in S\right) \leq \exp\left(\epsilon\right) \Pr\left(\mathcal{M}\left(z\right) \in S\right),$$

for any $y, z \in \mathfrak{R}$ and any $S \subset \mathfrak{S}$.

The idea is that the distributions produced by y and z are absolutely continuous to each other with the multiplier $\exp\left(\epsilon\right)$. With ϵ close to 0, these two distributions should look similar, and it will be quite difficult to infer whether it is y or z. When it comes to the rating vector, it should be $\Pr\left(\mathcal{M}\left(x^{(1)}\right) \in S\right) \leq \exp\left(\epsilon\right) \Pr\left(\mathcal{M}\left(x^{(2)}\right) \in S\right)$, where $x^{(1)}$ and $x^{(2)}$ are two different rating vectors, which may represent two different users. Hence, the outputs of all users are mixed up and thus indistinguishable.

2.2 Modified Laplace mechanism

In this subsection, we introduce modified Laplace mechanism. The name comes from Laplace mechanism [5], which works only on continuous metrizable space. In order to handle missing value, we will modify it a bit. For convenience, we suppose that the data are normalized into the interval $[-1, 1]$, and let the question mark ? denote the missing value.

Definition 2. For any $\epsilon \geq 0$, $\xi \sim \text{Laplace}(0, \frac{2}{\epsilon})$ is a random variable, and $\zeta \sim \text{Bernoulli}(\frac{\exp(\epsilon/2)}{\exp(\epsilon/2)+1})$ is a random variable independent to ξ. A modified Laplace mechanism $\mathcal{M}(x) = (\mathcal{M}(x_1), \mathcal{M}(x_2), \ldots, \mathcal{M}(x_n))$ is defined as

$$\mathcal{M}(x_i) = \begin{cases} \zeta \cdot (x_i + \xi) + (1-\zeta) \cdot ? & : x_i \in [-1, 1] \\ \zeta \cdot ? + (1-\zeta) \cdot \xi & : x_i = ? \end{cases}, \tag{1}$$

where by convention, $1 \cdot ? = ?$, $0 \cdot ? = 0$ and $? + 0 = ?$.

The idea is that besides adding Laplace noise, we randomly remove and create some ratings as well. This mechanism can be proven to be differentially private.

Theorem 1. *Modified Laplace mechanism is $n\epsilon$-differentially private.*

2.3 Randomized response mechanism

In this subsection, we present randomized response mechanism, which works on discrete data. Note that [6, 7] also use this term, and their definitions are even different between them. Our use of this term is closer to [7], and we will adapt it to rating data.

Recommender systems rarely allow users to give continuous ratings. Instead, they often ask the user to rate an item by one to five stars. These ratings are surely ordinal, but we just ignore the order of the set. Along with the missing rating, we consider them as cardinal numbers. In the following definition, the number 0 can be seen as the missing rating, and the numbers $1, 2, \ldots, d$ can be seen as the number of stars.

Definition 3. Let $W = \{0, 1, 2, ..., d\}$ be a set of finite cardinality. For any $\epsilon \geq 0$, any $i \in W$, ξ_i is a (independent) random variable with support on W whose probability mass function is defined by $p_i(j) = \frac{\exp(\epsilon)^{I(j=i)}}{\exp(\epsilon)+d}$, for any $j \in W$. A randomized response mechanism $\mathcal{M}(x) = (\mathcal{M}(x_1), \mathcal{M}(x_2), \ldots, \mathcal{M}(x_n))$ is defined as $\mathcal{M}(x_k) = \xi_{x_k}$.

The idea is that the transformed ratings (including missing ratings) will most likely remain the same as the original ratings, but there is still possibility that they are transformed to other ratings (with equal probability). One can prove that this mechanism is differentially private.

Theorem 2. *Randomized response mechanism is $n\epsilon$-differentially private.*

3 Utility

A natural question is that whether the transformed ratings are useful. If the transformed ratings produce nonsense, then there will be no meaning of this transform although the user privacy is protected. This question can be decomposed into two subquestions: what the usefulness means and what could be the possible way to make it useful. For the first question, we will use the statistical estimation framework. And for the second one, we will use the low-rank matrix completion method.

We start with the framework. There are m users and n items in the universe. $\Theta_{m \times n}$ is the unknown matrix of true ratings that each user will give to each item. This is a dense matrix without any missing values. However, since there are so many items, users are not able to test every item and their ratings are corrupted by noise. What we finally observe is a sparse matrix $X_{m \times n}$, which could be regarded as some approximation of Θ. Then, we apply either of our mechanism on X to generate the transformed rating matrix $Z_{m \times n}$. Since our mechanism is computed elementwisely, this can be

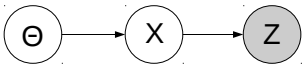

Figure 1: Rating generating process

done locally at each user's computer. After that, each user sends their transformed rating vector to the service provider, who observes the matrix Z. The service provider's goal is to recover Θ from Z.

Now we will show how it is possible to recover Θ from Z instead of from X. As mentioned in Section 1, there are quite a few methods. Interesting readers can refer to [1, 8]. Here we will only present one method, but the analysis can be generalized to all methods.

This method is low-rank matrix completion. We suppose Θ is a low-rank matrix, i.e. $\text{rank}(\Theta) = r \ll \min(m, n)$. If the true rating matrix is low-rank, then we are able to approximately recover it with a few corrupted ratings under certain condition such as *restricted isometry property* (RIP) [9].

Definition 4. Let Ω_Z denote the support of the non-missing ratings of Z. A projection operator $\mathcal{P}_{\Omega_Z}$ satisfies restricted isometry property if it obeys

$$(1-\alpha)\left\|A\right\|_F^2 \leq \tfrac{1}{p}\left\|\mathcal{P}_{\Omega_Z}(A)\right\|_F^2 \leq (1+\alpha)\left\|A\right\|_F^2, \tag{2}$$

for any matrix A with sufficiently small rank and $\alpha \in (0,1)$ sufficiently small, where p is the proportion of non-missing values of Z and $\left\|\cdot\right\|_F$ is the Frobenius norm.

The recover process is described as follow. Suppose that

$$\rho := \left\|\mathcal{P}_{\Omega_Z}(\Theta - Z)\right\|_F < \infty. \tag{3}$$

Our estimator $\hat{\Theta}$ is obtained from the following optimization problem

$$\begin{aligned} \arg\min_{M} \quad & \left\|M\right\|_* \\ \text{s.t.} \quad & \left\|\mathcal{P}_{\Omega_Z}(M - Z)\right\|_F \leq \rho, \end{aligned} \tag{4}$$

where $\left\|\cdot\right\|_*$ is the nuclear norm (a.k.a. trace norm).

Under low-rank hypothesis and RIP, [9] proved

$$\left\|\hat{\Theta} - \Theta\right\|_F \leq C_0 p^{-1/2}\rho \tag{5}$$

for some numerical constant C_0. This means that the estimation error on the whole matrix is proportional to the error on the support of the observed matrix, which means that the recover method enjoys a kind of stability against the noise quantified by ρ. Of course, this noise includes not only the noise intrinsic in the problem (i.e. between Θ and X) but also the noise artificially introduced through the mechanism (i.e. between X and Z). We see how easily the traditional analysis techniques can be seamlessly transplanted to the new setting.

What remains is just to give an upper bound of ρ. Let Ω_X denote the support of the non-missing ratings of X, and $s := |\Omega_X|$ be the number of non-missing ratings. Suppose that $\left\|\mathcal{P}_{\Omega_X}(\Theta - X)\right\|_F \leq \rho_0\sqrt{s} < \infty$, for some small constant ρ_0. This hypothesis is quite realistic. Indeed, this is what we need if we want to recover Θ from the untransformed ratings X. Then we have the following theorems.

Theorem 3. *Let $\gamma \in (0,1)$ be the level of tolerance. With probability at least $1-\gamma$, the Z generated by modified Laplace mechanism satisfies*

$$\rho \leq \rho_0\sqrt{s} + \frac{4}{\epsilon}\sqrt{\frac{s}{\gamma}} + \sqrt{\frac{2mn}{(e^{\frac{\epsilon}{2}}+1)\gamma}\left(1 + \frac{8}{\epsilon^2}\right)}; \tag{6}$$

with probability at least $1-\gamma$*, the* Z *generated by randomized response mechanism satisfies*

$$\rho \leq \rho_0\sqrt{s} + 2(d-1)\sqrt{\frac{2mnd}{(e^\epsilon + d)\gamma}}. \tag{7}$$

When the privacy parameter ϵ increases toward infinity, the above upper bounds decrease toward $\rho_0\sqrt{s}$. This means that the less the level of differential privacy is, the more accurate the data are and then the more precise estimation we will have. This is intuitive, since the larger ϵ is, the less extra noise we introduce into the data. In practice, it is desirable to choose an ϵ which makes the entire upper bound match the order of $\rho_0\sqrt{s}$. Combining these bounds with (5), we assure the utility of our transformed ratings.

4 Proof

4.1 Proof of Theorem 1

Proof. According to the values of (x, y) and S, we should divide it into nine cases. We only consider non empty set of S, since the empty set case is trivial.

i) $(x, y) \in [-1, 1]^2$ and $S \in \mathcal{R}([-1, 1])$

$$\begin{aligned}
\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} &= \frac{\Pr(\zeta_1 = 1, x + \xi_1 \in S)}{\Pr(\zeta_2 = 1, y + \xi_2 \in S)} \\
&= \frac{\Pr(\zeta_1 = 1)\Pr(x + \xi_1 \in S)}{\Pr(\zeta_2 = 1)\Pr(y + \xi_2 \in S)} \\
&= \frac{\Pr(x + \xi_1 \in S)}{\Pr(y + \xi_2 \in S)} \\
&\leq e^\epsilon.
\end{aligned}$$

ii) $(x, y) \in [-1, 1]^2$ and $S = \{?\}$

$$\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} = \frac{\Pr(\zeta_1 = 0)}{\Pr(\zeta_2 = 0)} = 1 \leq e^\epsilon.$$

iii) $(x, y) \in [-1, 1]^2$ and S is more than $\{?\}$

$$\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} = \frac{\Pr(\zeta_1 = 1, x + \xi_1 \in S) + \Pr(\zeta_1 = 0)}{\Pr(\zeta_2 = 1, y + \xi_2 \in S) + \Pr(\zeta_2 = 0)} \leq e^\epsilon.$$

iv) $x \in [-1, 1]$, $y =?$, and $S \in \mathcal{R}([-1, 1])$

$$\begin{aligned}
\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} &= \frac{\Pr(\zeta_1 = 1, x + \xi_1 \in S)}{\Pr(\zeta_2 = 0, \xi_2 \in S)} \\
&= \frac{\Pr(\zeta_1 = 1)\Pr(x + \xi_1 \in S)}{\Pr(\zeta_2 = 0)\Pr(\xi_2 \in S)} \\
&\leq e^{\frac{\epsilon}{2}} \cdot e^{\frac{\epsilon}{2}} \\
&= e^\epsilon.
\end{aligned}$$

v) $x \in [-1, 1]$, $y =?$, and $S = \{?\}$

$$\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} = \frac{\Pr(\zeta_1 = 0)}{\Pr(\zeta_2 = 1)} = e^{-\frac{\epsilon}{2}} \leq e^{\epsilon}.$$

vi) $x \in [-1, 1]$, $y =?$ and S is more than $\{?\}$

$$\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} = \frac{\Pr(\zeta_1 = 1, x + \xi_1 \in S) + \Pr(\zeta_1 = 0)}{\Pr(\zeta_2 = 0, \xi_2 \in S) + \Pr(\zeta_2 = 1)} \leq e^{\epsilon}.$$

vii) $y \in [-1, 1]$, $x =?$, and $S \in \mathcal{R}([-1, 1])$

$$\begin{aligned}\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} &= \frac{\Pr(\zeta_1 = 0, \xi_1 \in S)}{\Pr(\zeta_2 = 1, y + \xi_2 \in S)} \\ &= \frac{\Pr(\zeta_1 = 0) \Pr(\xi_1 \in S)}{\Pr(\zeta_2 = 1) \Pr(y + \xi_2 \in S)} \\ &\leq e^{-\frac{\epsilon}{2}} \cdot e^{\frac{\epsilon}{2}} \\ &\leq e^{\epsilon}.\end{aligned}$$

viii) $y \in [-1, 1]$, $x =?$, and $S = \{?\}$

$$\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} = \frac{\Pr(\zeta_1 = 1)}{\Pr(\zeta_2 = 0)} = e^{\frac{\epsilon}{2}} \leq e^{\epsilon}.$$

ix) $y \in [-1, 1]$, $x =?$ and S is more than $\{?\}$

$$\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} = \frac{\Pr(\zeta_1 = 0, \xi_1 \in S) + \Pr(\zeta_1 = 1)}{\Pr(\zeta_2 = 1, y + \xi_2 \in S) + \Pr(\zeta_2 = 0)} \leq e^{\epsilon}.$$

□

4.2 Proof of Theorem 2

Proof. For any $(x, y) \in W^2$, we have

$$\frac{\Pr(\mathcal{M}(x) \in S)}{\Pr(\mathcal{M}(y) \in S)} = \frac{\sum_{s \in S} \Pr(\mathcal{M}(x) = s)}{\sum_{s \in S} \Pr(\mathcal{M}(y) = s)}. \tag{8}$$

s could be three kinds of values: x, y and the other. So we divide it into three cases.

i) $s = x$

$$\frac{\Pr(\mathcal{M}(x) = x)}{\Pr(\mathcal{M}(y) = x)} = \frac{\frac{e^{\epsilon}}{e^{\epsilon}+d}}{\frac{1}{e^{\epsilon}+d}} = e^{\epsilon}.$$

ii) $s = y$

$$\frac{\Pr(\mathcal{M}(x) = y)}{\Pr(\mathcal{M}(y) = y)} = \frac{\frac{1}{e^{\epsilon}+d}}{\frac{e^{\epsilon}}{e^{\epsilon}+d}} = e^{-\epsilon} \leq e^{\epsilon}.$$

iii) $s \neq x$ and $s \neq y$

$$\frac{\Pr(\mathcal{M}(x)=s)}{\Pr(\mathcal{M}(y)=s)} = \frac{\frac{1}{e^\epsilon+d}}{\frac{1}{e^\epsilon+d}} = 1 \leq e^\epsilon.$$

So in either way, the fraction is no more than e^ϵ. Join these equations into (8), we get

$$\frac{\Pr(\mathcal{M}(x)\in S)}{\Pr(\mathcal{M}(y)\in S)} \leq \frac{\sum_{s\in S}\Pr(\mathcal{M}(y)=s)e^\epsilon}{\sum_{s\in S}\Pr(\mathcal{M}(y)=s)} = e^\epsilon.$$

□

4.3 Proof of Theorem 3

Proof. Firstly, we decompose 3 into three terms.

$$\begin{aligned}\left\|\mathcal{P}_{\Omega_Z}(\Theta - Z)\right\|_F &= \left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(\Theta - X) + \mathcal{P}_{\Omega_Z\cap\Omega_X}(X - Z) + \mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta - Z)\right\|_F \\ &\leq \left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(\Theta - X)\right\|_F + \left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X - Z)\right\|_F + \left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta - Z)\right\|_F.\end{aligned}$$

The first term

$$\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(\Theta - X)\right\|_F \leq \left\|\mathcal{P}_{\Omega_X}(\Theta - X)\right\|_F \leq \rho_0\sqrt{s}.$$

Then, we calculate the mathematical expectation of the square of the second term.

$$\begin{aligned}\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X - Z)\right\|_F^2 &= \mathbb{E}\left[\sum_{(i,j)\in\Omega_Z\cap\Omega_X}(X_{ij} - Z_{ij})^2\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\sum_{(i,j)\in\Omega_Z\cap\Omega_X}(X_{ij} - Z_{ij})^2|\Omega_X,\Omega_Z\right]\right] \\ &= \mathbb{E}\left[\sum_{(i,j)\in\Omega_Z\cap\Omega_X}\mathbb{E}\left[(X_{ij} - Z_{ij})^2|\Omega_X,\Omega_Z\right]\right].\end{aligned} \tag{9}$$

In the same way, we also calculate for the third term.

$$\begin{aligned}\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta - Z)\right\|_F^2 &= \mathbb{E}\left[\sum_{(i,j)\in\Omega_Z\setminus\Omega_X}(\Theta_{ij} - Z_{ij})^2\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\sum_{(i,j)\in\Omega_Z\setminus\Omega_X}(\Theta_{ij} - Z_{ij})^2|\Omega_X,\Omega_Z\right]\right] \\ &= \mathbb{E}\left[\sum_{(i,j)\in\Omega_Z\setminus\Omega_X}\mathbb{E}\left[(\Theta_{ij} - Z_{ij})^2|\Omega_X,\Omega_Z\right]\right].\end{aligned} \tag{10}$$

For modified Laplace mechanism

$$\begin{aligned}&\mathbb{E}\left[(X_{ij} - Z_{ij})^2|\Omega_X,\Omega_Z\right] = 2\left(\frac{2}{\epsilon}\right)^2 = \frac{8}{\epsilon^2}, \quad \forall i,j\in\Omega_X\cap\Omega_Z; \\ &\mathbb{E}\left[(\Theta_{ij} - Z_{ij})^2|\Omega_X,\Omega_Z\right] = \Theta_{ij}^2 + 2\left(\frac{2}{\epsilon}\right)^2 \leq 1 + \frac{8}{\epsilon^2}, \quad \forall i,j\in\Omega_Z\setminus\Omega_X.\end{aligned}$$

Join these into (9) and (10), we get

$$\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X-Z)\right\|_F^2 \le \frac{8s}{\epsilon^2}$$
$$\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta-Z)\right\|_F^2 \le \frac{mn}{e^{\frac{\epsilon}{2}}+1}\left(1+\frac{8}{\epsilon^2}\right),$$

where $s := |\Omega_X|$.

Then,

$$\Pr\left(\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X-Z)\right\|_F > \sqrt{\frac{16s}{\epsilon^2\gamma}}\right) \le \frac{\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X-Z)\right\|_F^2}{\frac{16s}{\epsilon^2\gamma}} = \frac{\gamma}{2}$$
$$\Pr\left(\left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta-Z)\right\|_F > \sqrt{\frac{2mn}{(e^{\frac{\epsilon}{2}}+1)\gamma}\left(1+\frac{8}{\epsilon^2}\right)}\right) \le \frac{\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta-Z)\right\|_F^2}{\frac{2mn}{(e^{\frac{\epsilon}{2}}+1)\gamma}\left(1+\frac{8}{\epsilon^2}\right)} = \frac{\gamma}{2}.$$

Finally, we have

$$\Pr\left(\left\|\mathcal{P}_{\Omega_Z}(\Theta-Z)\right\|_F \le \rho_0\sqrt{s} + \frac{4}{\epsilon}\sqrt{\frac{s}{\gamma}} + \sqrt{\frac{2mn}{(e^{\frac{\epsilon}{2}}+1)\gamma}\left(1+\frac{8}{\epsilon^2}\right)}\right) \ge 1-\gamma.$$

For randomized response mechanism

$$\mathbb{E}\left[(X_{ij}-Z_{ij})^2|\Omega_X,\Omega_Z\right] \le (d-1)^2\frac{d}{e^\epsilon+d}, \quad \forall i,j\in\Omega_X\cap\Omega_Z;$$
$$\mathbb{E}\left[(\Theta_{ij}-Z_{ij})^2|\Omega_X,\Omega_Z\right] = (d-1)^2, \quad \forall i,j\in\Omega_Z\setminus\Omega_X.$$

Join these into (9) and (10), we get

$$\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X-Z)\right\|_F^2 \le (d-1)^2\frac{sd}{e^\epsilon+d},$$
$$\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta-Z)\right\|_F^2 \le \frac{mnd}{e^\epsilon+d}(d-1)^2,$$

where $s := |\Omega_X|$.

Then,

$$\Pr\left(\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X-Z)\right\|_F > \sqrt{\frac{2sd(d-1)^2}{(e^\epsilon+d)\gamma}}\right) \le \frac{\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\cap\Omega_X}(X-Z)\right\|_F^2}{\frac{2sd(d-1)^2}{(e^\epsilon+d)\gamma}} = \frac{\gamma}{2}$$
$$\Pr\left(\left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta-Z)\right\|_F > \sqrt{\frac{2mnd(d-1)^2}{(e^\epsilon+d)\gamma}}\right) \le \frac{\mathbb{E}\left\|\mathcal{P}_{\Omega_Z\setminus\Omega_X}(\Theta-Z)\right\|_F^2}{\frac{2mnd(d-1)^2}{(e^\epsilon+d)\gamma}} = \frac{\gamma}{2}.$$

Finally, since $s \le mn$, we have

$$\Pr\left(\left\|\mathcal{P}_{\Omega_Z}(\Theta-Z)\right\|_F \le \rho_0\sqrt{s} + 2(d-1)\sqrt{\frac{2mnd}{(e^\epsilon+d)\gamma}}\right) \ge 1-\gamma.$$

□

References

[1] Francesco Ricci, Lior Rokach, and Bracha Shapira. *Introduction to recommender systems handbook*. Springer, 2011.

[2] James Bennett and Stan Lanning. The netflix prize. In *Proceedings of KDD cup and workshop*, volume 2007, page 35, 2007.

[3] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, pages 111–125. IEEE, 2008.

[4] Frank McSherry and Ilya Mironov. Differentially private recommender systems: building privacy into the net. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 627–636. ACM, 2009.

[5] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography*, pages 265–284. Springer, 2006.

[6] John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy, data processing inequalities, and statistical minimax rates. *arXiv preprint arXiv:1302.3203*, 2013.

[7] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. In *Advances in neural information processing systems*, pages 2879–2887, 2014.

[8] Charu C Aggarwal. *Recommender Systems: The Textbook*. Springer, 2016.

[9] M Fazel, E Candes, B Recht, and P Parrilo. Compressed sensing and robust recovery of low rank matrices. In *Signals, Systems and Computers, 2008 42nd Asilomar Conference on*, pages 1043–1047. IEEE, 2008.